

Distance Metric Ensemble Learning and the Andrews-Curtis Conjecture

Krzysztof Krawiec^{*1} and Jerry Swan^{†2}

¹Poznan University of Technology, Piotrowo 2, 60-965 Poznan, Poland.

²University of York, Deramore Lane, York, YO10 5GH, UK.

June 7, 2016

Abstract

Motivated by the search for a counterexample to the Poincaré conjecture in three and four dimensions, the Andrews-Curtis conjecture was proposed in 1965. It is now generally suspected that the Andrews-Curtis conjecture is false, but small potential counterexamples are not so numerous, and previous work has attempted to eliminate some via combinatorial search. Progress has however been limited, with the most successful approach (breadth-first-search using secondary storage) being neither scalable nor heuristically-informed. A previous empirical analysis of problem structure examined several heuristic measures of search progress and determined that none of them provided any useful guidance for search. In this article, we induce new quality measures directly from the problem structure and combine them to produce a more effective search driver via ensemble machine learning. By this means, we eliminate 19 potential counterexamples, the status of which had been unknown for some years.

Keywords: Andrews-Curtis conjecture; meta-heuristic search; machine learning.

1 Introduction

The Andrews-Curtis conjecture (ACC) [Andrews and Curtis 1965] dates back to 1965 and is an open problem of widespread interest in low-dimensional topology [Wright 1975; Hog-Angeloni and Metzler 1993] and combinatorial group theory [Burns and Macedońska 1993; Schupp and III 1999]. It originated in the search for a counterexample to the Poincaré conjecture in three and four dimensions. Subsequent to the proof of the Poincaré conjecture [Perelman 2003], it is generally suspected that ACC is false. Attention has therefore shifted to potential counterexamples to ACC, of which relatively few of likely computational tractability are known [Bridson 2006].

ACC can be stated in both group-theoretic and topological terms. We proceed via the elementary theory of group presentations [Johnson 1990]. A finite presentation $\langle g_1, \dots, g_m | r_1, \dots, r_n \rangle$ is said to be *balanced* if m is equal to n . The *trivial presentation of the trivial group* of rank r is the balanced presentation $\langle g_1, \dots, g_r | g_1, \dots, g_r \rangle$. For conciseness, we sometimes denote the inverse of a generator by capitalization, e.g. $B = b^{-1}$.

The group-theoretic version of ACC states that “every balanced presentation of the trivial group can be transformed into the trivial presentation via some sequence of *AC-moves*” [Burns and Macedońska 1993]. For relators r_i, r_j , the AC-moves are:

^{*}krawiec@cs.put.poznan.pl

[†]jerry.swan@york.ac.uk

1. **AC1.** $r_i \rightarrow r_i^{-1}$ (inversion of a relator)
2. **AC2.** $r_i \rightarrow r_i r_j, i \neq j$ (multiplication of one relator by another)
3. **AC3.** $r_i \rightarrow g^{\mp 1} r_i g^{\pm 1}$, (*conjugation* of a relator by some generator g)

We say that a sequence of AC-moves that connects a source presentation p to the trivial group is an *AC-trivialisation* of p . The contribution of this article is a novel approach to searching for AC-trivializations, leading to the elimination of 19 of the potential counterexamples described at the end of the next section. The proposed metaheuristic algorithm combines offline learning and online ensemble approaches Kittler et al. [1998].

2 Previous Work

It is possible to investigate potential counterexamples to ACC using combinatorial search techniques such as genetic algorithms [Holland 1992] and breadth-first search [Havas and Ramsay 2003; Bowman and McCaul 2006]. The search is therefore for some sequence of moves connecting the trivial group to a potential counterexample. Metaheuristic approaches are guided by a *fitness function*, an ordering on solution states that gives a heuristic measure of the quality of a solution. Despite both group-theoretic [Miasnikov and Myasnikov 2003] and metaheuristic [Miasnikov 1999] approaches, the state-of-the-art since 2003 has been breadth-first search [Havas and Ramsay 2003], subsequently extended to efficiently index secondary storage [Bowman and McCaul 2006]. More recently, [Lisitsa 2013] used the alternative approach of first-order theorem proving to obtain trivializations for all previously-eliminated potential counterexamples.

The AC-moves themselves form a group (denoted AC_n) under their action on balanced presentations of rank n : AC1 and AC3 are self-inverse and AC2 is inverted by multiplication by the inverse of the source relator. One can therefore either start from a potential counterexample and search ‘forwards’ towards the trivial presentation or else start at the trivial presentation and apply inverse moves. In this manner,

Havas and Ramsay make use of bidirectional breadth-first search [Havas and Ramsay 2003], terminating with success if the search frontiers intersect. For a balanced presentation of rank n , there are $3n^2$ AC-moves, and hence $(3n^2)^l$ move sequences of length l . By the group property of AC_n , it is clear that the effective length of many sequences is lower, e.g. the immediate re-application of self-inverse moves will always yield the previously-encountered presentation. In practice, for the rank 2 case, Havas and Ramsay note that the theoretical branching factor of 12 tends to average at around 8 in the low-depth unconstrained investigations they performed.

For breadth-first search, constraints on relator length are used to make the state space finite (and furthermore tractable), and ‘total length of relators’ has been used as an estimate of problem difficulty. In these terms, the smallest potential counterexample is $AK_3 = \langle a, b | a^3 B^4, abaBAB \rangle$ of length 13 due to Akbulut and Kirby [Akbulut and Kirby 1985]. In [Bowman and McCaul 2006] Bowman and McCaul exhaustively enumerated the constrained search space for AK_3 for maximum individual relator lengths from 10 to 17 inclusive, but were unable to find a solution sequence, despite enumerating 85 million presentations and taking 93 hours on an IBM z800 mainframe. It is clearly therefore necessary to explore alternative approaches. In this paper, we investigate the application of a more informed version of metaheuristic search than has previously been attempted. Metaheuristic search has an associated *fitness landscape* [Wright 1932; Stadler 1995], i.e. a graph in which the vertices are (potential) solutions and the edges of the graph represent the operations for transforming a solution into its neighbour. Previous work by Swan et al. [Swan et al. 2012] has explored alternatives to relator length as a fitness measure (e.g. edit distance) and determined that fitness does not correlate well with the distance (expressed in terms of number of edges traversed) to a solution [Jones and Forrest 1995].

In a broader context, with the exception of work by Spector et al. [Spector et al. 2008] which uses genetic programming [Koza 1992] to discover terms with specific properties in finite universal algebras, we are not aware of any significant applications of machine

learning techniques to algebraic problems of general interest.

3 Problem Instances

Determining if a balanced presentation actually represents the trivial group (and is therefore a potential counterexample) is a nontrivial task [Edjvet et al. 2001; Miasnikov and Myasnikov 2003] in its own right. A regularly-updated collection of balanced presentations arising from computational and algebraic investigations into irreducible cyclicly presented groups performed since 2001 [Edjvet 2003; Edjvet and Spanu 2011; Cremona and Edjvet 2010; Edjvet and Swan 2014] is maintained at [Edjvet 2013]. These fall into two categories:

- Presentations T_i known to be trivial. These are potential counterexamples to ACC, and are of interest as described above.
- Presentations O_i for which triviality is an open question. In such cases, obtaining an AC-trivialisation additionally provides answers to a question of longstanding interest due to Dunwoody [Dunwoody 1995].

These instances have resisted further investigation by both algebraic and computational approaches over a number of years. For the instances O_i , approaches have included application of string-rewriting systems via the automatic groups software packages KB-MAG [Holt 1995] and MAF [Williams 2010] and the computer algebra package MAGMA (via the algorithms for simple quotients or low-index subgroups [Bosma et al. 1997]).

4 Methodology

As discussed above, in combinatorial terms, algorithmic verification of AC counterexamples is a search problem, with states corresponding to presentations and neighborhood defined by the set of AC moves. No fitness function is known that would efficiently guide search in this space, and, as shown in [Swan et al.

2012], the functions used in past studies do not correlate well with the actual distance to the search target (i.e. the number of AC moves required to reach the trivial presentation). It is therefore unsurprising that the greatest successes to date have not been heuristically informed. One of the conclusions of Swan et al. was that an adaptive or penalty-driven fitness function may allow metaheuristic approaches to outperform breadth-first search.

The central claim of this study is that, in the likely absence of a unique, global and efficiently-computable fitness function, a potentially useful substitute for it can be *learned* from the problem. In the following, we refer to such substitute functions as *distance metrics*. A distance metric should exhibit (some degree of) correlation with the actual (in practice unknown) distance from the search target. In contrast to conventional fitness functions, we do not require it to be globally minimized at the search target. In outline, the method is split into three phases:

1. Preparation of a training set of presentations.
2. Offline learning of a set of distance metrics on presentations.
3. Online search for trivialising sequences of AC-moves, using genetic algorithms equipped with a fitness measure which is informed by an ensemble of the generated distance metrics.

Of these phases, detailed in subsequent sections, phases 2 and 3 are implemented as a generational genetic algorithm [Holland 1992].

4.1 Preparation of training set

This phase consists in generation of a training set of *fitness cases*, i.e., examples with which the distance metrics are trained. Each fitness case is a pair (p, l) , where p is a randomly-chosen presentation a small number of AC-moves from the trivial presentation, and l is the length of the shortest path that trivialises p (referred to as *distance* in the following). It is clearly not possible in general to obtain this path directly by starting at some arbitrary p , since this would be equivalent to showing that p is

AC-trivializable. This forces us to devise a different approach for drawing the presentations for fitness cases. We start from the trivial presentation t , perform a reverse random walk, and terminate it at a state p if walk length exceeds 60 or the total length of relators reaches 60. Then, we attempt to find the corresponding inverse forward walk from p to t . Since all AC-moves other than multiplication are self-inverse, the length of forward and reverse walks will generally correlate well, however a move such as $(AABaB, BAbaB) \rightarrow (AB, BAbaB)$ from the proof of the AK_2 example from [Havas and Ramsay 2003] requires several moves to invert.

In principle, we could perform the reverse random walk by simply applying a random sequence of AC-moves of a given length to the trivial presentation. Since such a walk is unlikely to be the shortest path, we instead explicitly build the graph of reverse moves rooted at the trivial presentation using breadth first search and subsequently sample walks from it.

The outcome of this stage is a set of fitness cases $T = \{(p, l)\}$. For all presentations of a given rank, it is sufficient to produce such a sample once, as all instances of potential ACC counterexamples dwell in the same search space.

4.2 Offline Distance Learning

Given a sample of fitness cases T , the goal of the next step is to learn an approximate unary distance metric $d : P \rightarrow \mathbb{R}^+$ that, for a given presentation p , predicts its distance from the trivial presentation t . Ideally, we would like to synthesize a function d such that $d(p) = l$ for every fitness case (p, l) in T (and possibly beyond it), but this cannot be done otherwise than by running a costly tree search from p and terminating once t has been reached. Rather than that, we attempt to learn *heuristic estimates* of l . Based on this motivation, we set our goal to learning d such that $d(p)$ and l are well *correlated*.

Technically, the process of learning the distance metric is realized as an evolutionary algorithm working with the population of *candidate solutions*, each of them representing a specific distance measure d . The fitness value of a given candidate solution d is calculated by applying d to all fitness cases in T and

computing the correlation coefficient between $d(p)$ and l . Depending on the setup, we employ linear correlation (Pearson) or rank-wise correlation (Kendall).

The heuristic distance estimates learned in this way are universal in representing domain knowledge that is common for all ACC instances of a given rank. This is another reason for treating this learning process as a separate stage (stage 2) of our workflow that precedes the actual solving of particular instances of ACC (Section 4.3). For the same reason, we refer to it as to *offline* distance learning.

The critical design choice concerns the representation of distance metrics. Previous studies resorted to total sum of relator lengths of p , edit distance between p and the trivial presentation t , or other generic metrics (see Section 2). Following [Swan et al. 2012], we posit that no significant correlation with the actual distance can be achieved without involving a greater degree of domain knowledge. On the other hand, manual design of metrics is time consuming, and likely to result in measures which suffer from unhelpful bias.

We therefore elected to represent the candidate metrics in the same manner as the solutions to the underlying ACC problem: each d is a sequence of ACC moves. When evaluated on a given presentation p , the moves in d are applied to p one by one, resulting in a certain presentation p' . The total relator length of the resulting presentation $|p'|$ is interpreted as the value of $d(p)$. The correlation of $d(p)$ with l forms the fitness of the candidate metrics.

The objective of this evolutionary distance learning is thus to synthesize a sequence of moves that ‘corrects’ the total relator length of a given presentation w.r.t. its actual distance from t (i.e., appropriately shortens and extends the relators). By resorting to correlation, we do not require the total relator length of the resulting presentation p' to be *equal* to the actual distance.

However, even with this relaxation, it would be naïve to assume that a metric with perfect correlation (over the set of all starting presentations p available in T) can be expressed using AC moves. Thus, rather than searching for a single ideal metric, we run evolutionary search 50 times and collect the best-of-run candidate metrics from all runs, forming so a *sam-*

ple of metrics D . The distance metrics obtained in particular runs are gathered in a set D that is a parameter of the subsequent step. This is an example of ensemble learning [Kittler et al. 1998], in which the deficiencies of inaccurate predictors are improved by generating a diverse collection of them and aggregating their outputs. The resulting ensemble is then expected to have greater accuracy than an individual predictor.

4.3 Online search for AC-trivialisations

The set of distance metrics D learned in the previous section allows us to devise fitness functions to guide the actual search for trivializations. In contrast to the previous two steps, this stage proceeds *online*, i.e., for each presentation (problem instance) independently.

Metaheuristic search is parameterized by three essential components, viz. a solution representation, a set of operators for changing or recombining solution representations, and a fitness measure. We adopt the same formulation for the first two of these as the genetic algorithms approach of Miasnikov [Miasnikov 1999], i.e. solutions are represented as sequences of AC-moves and operators are the insertion, deletion and substitution of a move. These operators take one solution as an argument and are by this token known as *mutations* in the terminology of genetic algorithms. No binary (two-argument) *crossover* search operators are applied in our setup.

Our fitness measure is based on the approximate metrics learned in the process described in Section 4.2. We consider two ways of conducting the selection process based on the set of metrics D learned there.

Single objective. In this variant, we apply the metrics in D to the training sample T of fitness cases and perform multiple linear regression of the obtained values against the reverse walk length l . In other words, a vector \mathbf{w} of weights w_i is found such that the linear combination of the distances

$$f(p) = \sum_{d_i \in D} w_i d_i(p) \quad (1)$$

minimizes the square root error with respect to l , i.e.

$$\min_{\mathbf{w}} \sum_{(p,l) \in T} (f(p) - l)^2.$$

The function $f(p)$ constructed in this way becomes the fitness that drives a conventional single-objective search as in the approach of Miasnikov [Miasnikov 1999].

Multi-objective. Recent work in evolutionary computation indicates that heuristic search can be more effective when driven with multiple objectives (fitness functions) rather than one [Jensen 2004]. Simultaneously maximizing multiple objectives that express various characteristics of candidate solutions is a natural means for maintaining population diversity and reduces the risk of premature convergence, i.e. all candidate solutions in the population becoming very similar to each other (which hinders explorations of the search space).

Following these observations, in the second variant we do not combine the particular metrics $d_i \in D$ into a common fitness as in (1), but treat every $d_i \in D$ as a separate objective. In the selection stage of evolutionary run, we use Non-dominated Sorting Genetic Algorithm (NSGA-II, [Deb et al. 2002]). Given a population, NSGA-II builds a Pareto-ranking based on dominance relation that spans the objectives, and then employs tournament selection on Pareto-ranks to select the solutions. Given two solutions with the same Pareto-rank, it prefers the one from the less ‘crowded’ part of Pareto-front.

It is a known that multiobjective selection methods like NSGA-II tend to become ineffective when the number of objectives is high. Given the 50 objectives gathered in D , it is very unlikely for any candidate solution (move sequence) to dominate on these objectives any other move sequence in a working population. In order to reduce the number of objectives used in the multiobjective variant, we employ a heuristic procedure that trims D to the 5 least correlated objectives, where correlation is calculated in the same manner as in Section 4.2, i.e. with respect to the sample of presentations prepared in Section 4.1.

In both single- and multi-objective scenarios, we employ settings which are quite conventional for evo-

lutionary algorithms. The initial population of size 1000 is seeded with random sequences of length 8. In each iteration (generation), tournament selection with tournament of size 7 is applied to appoint the ‘parent’ candidate solutions that are then modified by search operators. In the single-objective variant, the selection is based on the scalar objective, while in the multi-objective variant it works with the ranks in the Pareto ranking induced by the dominance relation.

The selected solutions undergo one of three possible modifications (*search operators*), with the accompanying probabilities:

- Insertion of a randomly selected AC-move at a random location of a sequence (prob. 0.1).
- Replacement of a move at random location with a randomly generated AC-move (prob. 0.8).
- Deletion of a move at random location (prob. 0.1).

Thanks to equal probability of insertion and deletion, the expected change of length of this suite of operators is zero. Nevertheless, preliminary experiments showed that longer sequences tend to obtain better fitness. Therefore, to prevent excessive growth, sequences longer than 70 moves are assigned the worst possible fitness (which almost always results in eliminating them from a population). On the other hand, to avoid wasting time on considering very short sequences that are unlikely to trivialize the presentation in question, we penalize sequences shorter than 8 moves in the same way. Finally, the same penalization is applied to the sequences that traverse presentations with the total relator length greater or equal to 200.

Search proceeded until a trivializing sequence was found, or the number of generations reached 100,000, or three-hour runtime elapsed, whatever came first. As the evolutionary search is stochastic by nature and depends on the choice of initial population, we repeated the runs for every presentation for 20 seeds of a random number generator. The entire experiment involved at least 10,000 evolutionary runs in total. The computations were conducted on a cluster

of workstations equipped with 4-core CPUs, running under the Simple Linux Utility for Resource Management (SLURM) software framework [Jette et al. 2002].

With regard to the balanced presentations upon which AC move sequences act, we adopt two canonicalization constraints that differ slightly from those used in [Bowman and McCaul 2006], defined as follows:

- **C1.** Relators are sorted in shortlex, i.e. ‘length then lexicographic’ order.
- **C2.** Relators are chosen to be the least representative (under shortlex ordering) modulo cyclic permutation and inversion, subject to the constraint that it is freely reduced. This weaker constraint is necessary since we cannot enforce cyclic reduction: to do so would obviate the AC3 conjugate moves.

These constraints reduce the size of the search space in the graph- and walk- generation phases, albeit at additional computational expense in sorting and determining equivalence.

5 Results

We conducted an extensive series of computational experiments on the T_i and O_i presentations introduced at the end of Section 2, using several variants of the workflow presented in Section 4. Table 1 presents the list of presentations that have been solved by this setting, i.e., demonstrated to be AC-trivializable.

As can be seen from Table 1, the obtained AC-trivialization sequences vary in length from 6 to 25. Their lengths prevent them from being presented here in full, so here we list here only the sequences for presentations T1 and T13 in Fig. 1, with the others available online¹. For brevity, we relabel as follows: $x_0 \mapsto a, X_0 \mapsto A, x_1 \mapsto b, X_1 \mapsto B$.

It is interesting to note that none of the rank 3 presentations and none of the ‘O’ instances (i.e. those of unknown triviality) were solved by this approach.

¹<http://www.cs.put.poznan.pl/kkrawiec/wiki/?n=Site.AndrewsCurtis>

Table 1: List of presentations solved (AC-trivialised) by the proposed approach.

Identifier	Presentation	Trivialization Length
T1	$\langle x_0, x_1 x_0^2 x_1 X_0 X_1, x_1^2 x_0 X_1 X_0 \rangle$	6
T5	$\langle x_0, x_1 x_0^2 x_1^2 X_0 X_1^2, x_1^2 x_0^2 X_1 X_0^2 \rangle$	10
T11	$\langle x_0, x_1 x_0^3 x_1^2 X_0^2 X_1^2, x_1^3 x_0^2 X_1^2 X_0^2 \rangle$	14
T13	$\langle x_0, x_1 x_0^2 x_1 X_0 x_1 X_0 X_1, x_1^2 x_0 X_1 x_0 X_1 X_0 \rangle$	7
T29	$\langle x_0, x_1 x_0^3 x_1^3 X_0^2 X_1^3, x_1^3 x_0^3 X_1^2 X_0^3 \rangle$	21
T31	$\langle x_0, x_1 x_0^3 x_1 X_0 x_1 X_0 X_1^2, x_1^3 x_0 X_1 x_0 X_1 X_0^2 \rangle$	10
T34	$\langle x_0, x_1 x_0^2 x_1^2 x_0 X_1 X_0^2 X_1, x_1^2 x_0^2 x_1 X_0 X_1^2 X_0 \rangle$	10
T35	$\langle x_0, x_1 x_0^2 x_1^2 X_0 x_1 X_0 X_1^2, x_1^2 x_0^2 X_1 x_0 X_1 X_0^2 \rangle$	24
T39	$\langle x_0, x_1 x_0^2 x_1 X_0 x_1^2 X_0 X_1^2, x_1^2 x_0 X_1 x_0^2 X_1 X_0^2 \rangle$	10
T56	$\langle x_0, x_1 x_0^4 x_1^3 X_0^3 X_1^3, x_1^4 x_0^3 X_1^3 X_0^3 \rangle$	25
T61	$\langle x_0, x_1 x_0^3 x_1^2 X_0 x_1 X_0^2 X_1^2, x_1^3 x_0^2 X_1 x_0 X_1^2 X_0^2 \rangle$	14
T63	$\langle x_0, x_1 x_0^3 x_1^2 X_0 X_1^3 X_0 x_1, x_1^3 x_0^2 X_1 X_0^3 X_1 x_0 \rangle$	24
T66	$\langle x_0, x_1 x_0^3 x_1 X_0^2 x_1^2 X_0 X_1^2, x_1^3 x_0 X_1^2 x_0^2 X_1 X_0^2 \rangle$	14
T67	$\langle x_0, x_1 x_0^3 x_1 X_0 x_1^2 X_0 X_1^3, x_1^3 x_0 X_1 x_0^2 X_1 X_0^3 \rangle$	22
T76	$\langle x_0, x_1 x_0^2 x_1 x_0 x_1 X_0 X_1 X_0 X_1, x_1^2 x_0 x_1 x_0 X_1 X_0 X_1 X_0 \rangle$	10
T81	$\langle x_0, x_1 x_0^2 x_1 X_0 x_1 X_0 X_1 x_0 X_1, x_1^2 x_0 X_1 x_0 X_1 X_0 x_1 X_0 \rangle$	19
T82	$\langle x_0, x_1 x_0^2 x_1 X_0 X_1 x_0 x_1 X_0 X_1, x_1^2 x_0 X_1 X_0 x_1 x_0 X_1 X_0 \rangle$	10
T84	$\langle x_0, x_1 x_0^2 X_1 x_0 x_1 X_0 x_1 X_0 X_1, x_1^2 X_0 x_1 x_0 X_1 x_0 X_1 X_0 \rangle$	15
T85	$\langle x_0, x_1 x_0 x_1 x_0 x_1 X_0^2 X_1 x_0 X_1, x_1 x_0 x_1 x_0 X_1^2 X_0 x_1 X_0 \rangle$	24

T1:

$$\begin{aligned}
&\langle a, b|a^2bAB, b^2aBA \rangle \xrightarrow{(b^2aBA)^A} \langle a, b|a^2bAB, ab^2aBA^2 \rangle \\
&\langle a, b|a^2bAB, ab^2aBA^2 \rangle \xrightarrow{ab^2aBA^2 * = a^2bAB} \langle a, b|ab, a^2bAB \rangle \\
&\langle a, b|ab, a^2bAB \rangle \xrightarrow{(a^2bAB)^b} \langle a, b|ab, Ba^2bA \rangle \\
&\langle a, b|ab, Ba^2bA \rangle \xrightarrow{(ab)^A} \langle a, b|a^2bA, Ba^2bA \rangle \\
&\langle a, b|a^2bA, Ba^2bA \rangle \xrightarrow{(a^2bA)^{-1}} \langle a, b|aBA^2, Ba^2bA \rangle \\
&\langle a, b|aBA^2, Ba^2bA \rangle \xrightarrow{Ba^2bA * = aBA^2} \langle a, b|B, aBA^2 \rangle
\end{aligned}$$

T13:

$$\begin{aligned}
&\langle a, b|a^2bAbAB, b^2aBaBA \rangle \xrightarrow{(b^2aBaBA)^A} \langle a, b|a^2bAbAB, ab^2aBaBA^2 \rangle \\
&\langle a, b|a^2bAbAB, ab^2aBaBA^2 \rangle \xrightarrow{ab^2aBaBA^2 * = a^2bAbAB} \langle a, b|ab, a^2bAbAB \rangle \\
&\langle a, b|ab, a^2bAbAB \rangle \xrightarrow{(ab)^A} \langle a, b|a^2bA, a^2bAbAB \rangle \\
&\langle a, b|a^2bA, a^2bAbAB \rangle \xrightarrow{(a^2bA)^{-1}} \langle a, b|aBA^2, a^2bAbAB \rangle \\
&\langle a, b|aBA^2, a^2bAbAB \rangle \xrightarrow{(aBA^2)^b} \langle a, b|BaBA^2b, a^2bAbAB \rangle \\
&\langle a, b|BaBA^2b, a^2bAbAB \rangle \xrightarrow{(a^2bAbAB)^b} \langle a, b|BaBA^2b, Ba^2bAbA \rangle \\
&\langle a, b|BaBA^2b, Ba^2bAbA \rangle \xrightarrow{BaBA^2b * = Ba^2bAbA} \langle a, b|A, Ba^2bAbA \rangle
\end{aligned}$$

Figure 1: The sequences of trivializing moves found for T1 and T13.

Examining elementary differences between presentations does not provide any very helpful guidance: for example, the Hamming distance H between (the first relators of) successfully solved presentations T81 and T82 is 4, whereas the distance between T81 and the unsolved T83 is only 2. If one takes generators as being equivalent to their inverses, then both $H(T81, T82)$ and $H(T81, T83)$ are zero.

As observed by Havas and Ramsay [Havas and Ramsay 2003], relator length behaves highly nonmonotonically along the path to a solution. In general, the highly discontinuous effect of free reduction on words means that it is difficult to discern any distinguishing characteristics of the successful trivialization sequences. One might speculate that one of the main reasons that the ACC remains unsolved is that, considered in terms of algorithmic information theory [Chaitin 1996], AC-trivializations are ‘nearly incompressible’, i.e. cannot be readily expressed by a function of significantly lower complexity than the sequence itself. Pending deeper algebraic insights, this apparent lack of ‘obviously exploitable’ structure lends further support for the learned bias of our approach.

6 Conclusion

The Andrews-Curtis conjecture is a longstanding open problem of interest to topologists and group theorists [Andrews and Curtis 1965]. Attempts to eliminate potential counterexamples to the conjecture via combinatorial search has seen no practical improvement since the exhaustive enumerative approach of [Bowman and McCaul 2006] in 2006. Informed by previous work that analysed fitness correlations in the associated fitness landscape [Swan et al. 2012], we generate new predictors of search progress by performing offline learning to obtain good fitness functions. These predictors take the form of random walks in the search space that are good correlates for a more naïve measure of solution quality (i.e. total length of relators). This is supplemented with an online approach that randomly samples a subset of predictors. By this means, we successfully solved 19

problem instances which have withstood human and machine approaches since 2001.

Many solutions obtained by this approach comprise 20 or more moves and are thus substantially longer than the ones systematically enumerated in [Bowman and McCaul 2006] (up to length 17). Assuming the effective branching factor of 8 (Section 2), a sequence of length 20 corresponds to a search tree of $8^{21} - 1 \approx 9.22 \times 10^{18}$ nodes, arguably much too large to be systematically searched using algorithms like breadth-first search with currently available computational resources. For problem T56, with the longest trivializing sequence found in this study (25 moves), the tree is still greater by five orders of magnitude (3×10^{23} nodes). Given the absence of universal fitness functions to efficiently guide the search [Swan et al. 2012] reliance on some form of machine learning appears essential in obtaining further solutions combinatorially.

Acknowledgments

K. Krawiec acknowledges support from the National Science Centre (Narodowe Centrum Nauki) grant number 2014/15/B/ST6/05205.

References

- S. Akbulut and R. Kirby. A potential smooth counterexample in dimension 4 to the Poincaré conjecture, the Schoenflies conjecture, and the Andrews Curtis conjecture. *Topology*, 24:375–390, 1985.
- J. J. Andrews and M. L. Curtis. Free groups and handlebodies. *Proceedings of the American Mathematical Society*, 16(2):192–195, 1965. ISSN 00029939.
- W. Bosma, J. Cannon, and C. Playoust. The MAGMA algebra system 1. The user language. *J. Symb. Comput.*, 24(3-4):235–265, 1997. ISSN 0747-7171.
- R. S. Bowman and S. B. McCaul. Fast searching for Andrews-Curtis trivializations. *Experimental Mathematics*, 15(3), 2006.

- M. Bridson. On the complexity of balanced presentations and the Andrews-Curtis conjecture. Preprint., 2006.
- R. G. Burns and O. Macedońska. Balanced presentations of the trivial group. *Bull. London Math. Soc.*, 25(6):513526, 1993.
- G. Chaitin. A new version of algorithmic information theory. *Complexity*, 1(4):55–59, 1996. ISSN 1099-0526. doi: 10.1002/cplx.6130010410. URL <http://dx.doi.org/10.1002/cplx.6130010410>.
- J. E. Cremona and M. Edjvet. Cyclically presented groups and resultants. *International Journal of Algebra and Computation*, 20(03):417–435, 2010.
- K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Evolutionary Computation, IEEE Transactions on*, 6(2):182–197, apr 2002. ISSN 1089-778X. doi: 10.1109/4235.996017.
- M. J. Dunwoody. Cyclic presentations and 3-manifolds. In A. C. Kim and D. L. Johnson, editors, *GroupsKorea 94 Proceedings*, page 4755, Berlin, 1995. de Gruyter.
- M. Edjvet. On irreducible cyclic presentations. *Journal of Group Theory*, 2003.
- M. Edjvet. Irreducible cyclicly presented groups. Available online at <https://goo.gl/8p9B2S>, 2013.
- M. Edjvet and B. Spanu. On a certain class of cyclically presented groups. *Journal of Algebra*, 346(1): 165 – 179, 2011.
- M. Edjvet and J. Swan. On irreducible cyclic presentations of the trivial group. *Experimental Mathematics*, 23(2):181–189, 2014. doi: 10.1080/10586458.2014.888379.
- M. Edjvet, P. Hammond, and N. Thomas. Cyclic presentations of the trivial group. *Experimental Mathematics*, 10:303306, 2001.
- G. Havas and C. Ramsay. Breadth-first search and the Andrews-Curtis conjecture. *International Journal of Algebra and Computation*, 13(1):61–68, 2003.
- C. Hog-Angeloni and W. Metzler. The Andrews-Curtis conjecture and its generalizations. *Two-dimensional homotopy and combinatorial group theory*, 197:365–380, 1993.
- J. H. Holland. *Adaptation in Natural and Artificial Systems*. MIT Press, Cambridge, MA, USA, 1992. ISBN 0-262-58111-6.
- D. Holt. KBMAG - Knuth-Bendix in Monoids and Groups. (software package and documentation available under the GAP algebra system as <http://www.gap-system.org/Packages/kbmag.html>), 1995.
- M. T. Jensen. Helper-objectives: Using multi-objective evolutionary algorithms for single-objective optimisation. *J. Math. Model. Algorithms*, 3(4):323–347, 2004. doi: 10.1007/s10852-005-2582-2. URL <http://dx.doi.org/10.1007/s10852-005-2582-2>.
- M. A. Jette, A. B. Yoo, and M. Grondona. SLURM: Simple Linux Utility for Resource Management. In *In Lecture Notes in Computer Science: Proceedings of Job Scheduling Strategies for Parallel Processing (JSSPP) 2003*, pages 44–60. Springer-Verlag, 2002.
- D. Johnson. *Presentations of Groups*, volume 15 of *London Math. Soc. Stud. Texts*. Cambridge University Press, Cambridge, 1990.
- T. Jones and S. Forrest. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In *Proceedings of the 6th International Conference on Genetic Algorithms*, pages 184–192, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc. ISBN 1-55860-370-0.
- J. Kittler, M. Hatef, R. Duin, and J. Matas. On combining classifiers. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(3): 226–239, Mar 1998. ISSN 0162-8828. doi: 10.1109/34.667881.
- J. R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992. ISBN 0-262-11170-5.

- A. Lisitsa. First-order theorem proving in the exploration of Andrews-Curtis conjecture. *TinyToCS*, 2, 2013. URL <http://tinytocs.org/vol2/papers/tinytocs2-lisitsa.pdf>.
- A. D. Miasnikov. Genetic algorithms and the Andrews-Curtis conjecture. *IJAC*, 9:671–686, 1999.
- A. D. Miasnikov and A. G. Myasnikov. Balanced presentations of the trivial group on two generators and the Andrews-Curtis conjecture. In W. Kantor and A. Seress, editors, *Groups and Computation III*, volume 23, 257–263, Berlin, 2003.
- G. Perelman. Ricci flow with surgery on three-manifolds. <http://arxiv.org/abs/math/0303109>, 2003.
- P. Schupp and M. C. III. Some presentations of the trivial group. In R. Gilman, editor, *Groups, Languages and Automata*, pages 113–115. American Mathematical Society, Contemporary Mathematics, Vol. 250, 1999.
- L. Spector, D. M. Clark, I. Lindsay, B. Barr, and J. Klein. Genetic programming for finite algebras. In M. K. et al, editor, *GECCO '08: Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pages 1291–1298, Atlanta, GA, USA, 12–16 July 2008. ACM. doi: doi:10.1145/1389095.1389343. URL <http://www.cs.bham.ac.uk/~wbl/biblio/gecco2008/docs/p1291.pdf>.
- P. F. Stadler. Landscapes and their correlation functions. Working papers, Santa Fe Institute, 1995.
- J. Swan, G. Ochoa, G. Kendall, and M. Edjvet. Fitness landscapes and the Andrews-Curtis conjecture. *International Journal of Algebra and Computation*, 22(02), 2012.
- A. Williams. Monoid Automata Factory, version 2.0.3. Software package available from <http://sourceforge.net/projects/maffsa/>, 2010.
- P. Wright. Group presentations and formal deformations. *Trans. Amer. Math. Soc.*, 208:161–169, 1975.
- S. Wright. The roles of mutation, inbreeding, cross-breeding and selection in evolution. In D. F. Jones, editor, *Proceedings of the Sixth International Congress on Genetics*, volume 1, pages 356–366, 1932.